Wen Wei Loh*, Beatrijs Moerkerke, Tom Loeys and Stijn Vansteelandt

Heterogeneous indirect effects for multiple mediators using interventional effect models

<https://doi.org/10.1515/em-2020-0023>

Received July 15, 2020; accepted December 8, 2020; published online January 2, 2021

Abstract: Decomposing an exposure effect on an outcome into separate natural indirect effects through multiple mediators requires strict assumptions, such as correctly postulating the causal structure of the mediators, and no unmeasured confounding among the mediators. In contrast, interventional indirect effects for multiple mediators can be identified even when – as often – the mediators either have an unknown causal structure, or share unmeasured common causes, or both. Existing estimation methods for interventional indirect effects require calculating each distinct indirect effect in turn. This can quickly become unwieldy or unfeasible, especially when investigating indirect effect measures that may be modified by observed baseline characteristics. In this article, we introduce simplified estimation procedures for such heterogeneous interventional indirect effects using interventional effect models. Interventional effect models are a class of marginal structural models that encode the interventional indirect effects as causal model parameters, thus readily permitting effect modification by baseline covariates using (statistical) interaction terms. The mediators and outcome can be continuous or noncontinuous. We propose two estimation procedures: one using inverse weighting by the counterfactual mediator density or mass functions, and another using Monte Carlo integration. The former has the advantage of not requiring an outcome model, but is susceptible to finite sample biases due to highly variable weights. The latter has the advantage of consistent estimation under a correctly specified (parametric) outcome model, but is susceptible to biases due to extrapolation. The estimators are illustrated using publicly available data assessing whether the indirect effects of self-efficacy on fatigue via self-reported post-traumatic stress disorder symptoms vary across different levels of negative coping among health care workers during the COVID-19 outbreak.

Keywords: conditional effects; direct and indirect effects; effect decomposition; effect modification; natural effect models.

Introduction

Mediation analysis is widely used to assess the effect of an exposure or treatment (A) on an outcome (Y) that is transmitted via an intermediate variable that lies on a causal pathway from A to Y . A formal framework for mediation analysis has been developed using counterfactual-based distribution-free definitions of *natural direct and indirect effects* (Pearl 2001; Robins and Greenland 1992). These developments elucidate the ignorability assumptions needed to identify the natural (in)direct effects, and explicates decomposing the total effect into a direct and an indirect effect regardless of the statistical model. Nonlinear models for the mediator and the outcome may therefore be used in practice under this framework. Notwithstanding these advances, natural (in)direct effects may be uninformative of real-life interventions as it may be unfeasible to set multiple variables simultaneously to individual-specific counterfactual values (Didelez, Dawid, and Geneletti 2006; Naimi, Kaufman, and MacLehose 2014; Petersen, Sinisi, and van der Laan 2006), and it is impossible to

*Corresponding author: Wen Wei Loh, Department of Data Analysis, Ghent University, Gent, Belgium,

E-mail: WenWei.Loh@UGent.be

Beatrijs Moerkerke and Tom Loeys, Department of Data Analysis, Ghent University, Gent, Belgium

Stijn Vansteelandt, Department of Applied Mathematics, Computer Science and Statistics, Ghent University, Ghent, Belgium;
Department of Medical Statistics, London School of Hygiene and Tropical Medicine, London, UK

perform experiments in which the identification assumptions for natural (in)direct effects are guaranteed to be satisfied (Robins and Richardson 2010).

The problem lies in that natural (in)direct effects are defined in terms of so-called cross-world counterfactuals (Robins and Richardson 2010) involving composite, or nested, counterfactuals for the mediator and outcome. Identification thus demands either specifying Non-Parametric Structural Equations Model (Pearl 2009) structures for all observed variables, or assuming independence of mediator and outcome counterfactuals under different exposures; see Andrews and Didelez (2020) for a discussion of the latter assumption. In particular, this “cross-world” independence assumption is violated when there are multiple mediators with one mediator affecting another mediator so that the former is an exposure-induced (also termed time-varying) confounder of the mediator-outcome relation for the latter (Avin, Shpitser, and Pearl 2005). Mediation analyses in most substantive applications involve multiple mediators, either because scientific interest is in investigating the effects transmitted via multiple candidate or putative mediators, or because certain confounders affected by exposure are concurrently perceived as competing mediators. Extensions of mediation analysis using natural effects to settings that involve multiple mediators are therefore restricted to (fine-grained) decompositions that assume (i) the mediators to either be independent (Lange, Rasmussen, and Thygesen 2013; Taguri, Featherstone, and Cheng 2018) or conform to a correctly postulated causal structure whereby certain path-specific effects are identifiable under further assumptions (Albert et al. 2019; Daniel et al. 2015; Steen et al. 2017; Vansteelandt and VanderWeele 2012), and (ii) the mediators share no hidden (or unmeasured) common causes. But in most realistic scenarios, the causal structure among the mediators and the hidden common causes of the mediators are unknown, thus violating the assumptions needed for identification.

In contrast, *interventional (in)direct effects*, first introduced by Geneletti (2007), Didelez, Dawid, and Geneletti (2006) and VanderWeele, Vansteelandt, and Robins (2014) for a single mediator, are not defined in terms of cross-world counterfactuals. Unlike natural effects that are defined in terms of individual-level (deterministic) interventions on the mediator, interventional effects consider population-level (stochastic) interventions that set the value of the mediator to a random draw from its counterfactual distribution. This distinction allows identification of interventional indirect effects in the context of multiple mediators under empirically verifiable assumptions that could be ensured in a randomized trial (Moreno-Betancur and Carlin 2018). Furthermore, the interpretation of interventional effects remains meaningful even when the exposure cannot be manipulated at the individual level, and the natural and interventional effects may coincide empirically (Micali et al. 2018). For example, Jackson and VanderWeele (2018) describe interventional (in)direct effects using race as the exposure and socioeconomic status as the mediator, without having to define nested potential outcomes with respect to race or the exposure effects of race. Vansteelandt and Daniel (2017) generalized the definition of interventional effects to the multiple mediator setting, and demonstrated that the joint indirect effect of an exposure on an outcome can be decomposed into separate indirect effects via each distinct mediator, and an indirect effect via the mediators’ mutual dependence, regardless of the underlying causal structure of the mediators. In particular, an interventional indirect effect via a mediator is defined to capture the combined effect along all underlying causal pathways leading from exposure (possibly via other mediators) to the mediator of interest, then from the mediator directly to the outcome. Interventional (in)direct effects for multiple mediators are well-defined and can be identified, and thus unbiasedly estimated, even when the directions of the causal effects between the mediators are unknown, or the mediators are manifestations of an underlying latent process, or the mediators share hidden common causes. Recent discussions of interventional (in)direct effects include comparisons with natural (in)direct effects (Lok 2019; Moreno-Betancur and Carlin 2018; Nguyen, Schmid, and Stuart 2019), longitudinal mediation (VanderWeele and Tchetgen Tchetgen 2017), and path-specific interpretations for multiple mediators (Lin and VanderWeele 2017).

In this article, we consider heterogeneous interventional (in)direct effects for multiple mediators that can vary across different levels of the observed baseline characteristics. We build on the framework of Vansteelandt and Daniel (2017) to develop interventional indirect effect measures that may be modified by the baseline covariates. The mediators and outcome can be continuous or noncontinuous. Existing estimation methods demand either deriving closed form expressions of the estimators when the assumed mean models for

the mediators and the outcome are linear, or separately calculating each (in)direct effect in turn when the assumed models are nonlinear. When estimating interventional indirect effects that may take on different values for different levels of the observed covariate(s), either the derivations become complicated and prone to miscalculations, or the calculations have to be repeated for each unique covariate level in turn. Existing methods for inferring heterogeneous interventional (in)direct effects can potentially be unwieldy or unfeasible in such settings. To address this shortcoming, we propose using *interventional effect models* for straightforward and simultaneous modeling of the interventional (in)direct effects when there are multiple mediators. Interventional effect models adopt the functional form of *natural effect models* for multiple mediators (Lange, Rasmussen, and Thygesen 2013; Steen et al. 2017), a class of marginal structural mean models (Robins 2000), that express the mean potential outcomes in terms of hypothetical exposure levels for each mediator. Indirect and direct effects can be readily encoded as causal parameters that index these models, while allowing for effects to vary between different covariate levels via interactions between the exposure levels and the covariates. Directly modeling the interventional (in)direct effects of interest can therefore simplify estimation and inference. We propose two estimation procedures for fitting interventional effect models for multiple mediators: inverse weighting by the counterfactual mediator density or mass functions, or Monte Carlo integration. The key advantage of the former is that no outcome model needs to be (correctly) specified. However, this estimator can be susceptible to non-negligible finite sample biases due to highly variable weights. In contrast, the latter estimator has the advantage of consistent estimation under a correctly assumed (parametric) model for the outcome. But specifying an outcome model that is coherent with the posited effect model can be difficult or impossible for nonlinear effect models due to noncollapsibility (Greenland, Robins, and Pearl 1999). This renders the latter estimator susceptible to biases due to extrapolation.

The remainder of this article is as follows. Notation is first introduced, and the heterogeneous interventional (in)direct effects are defined following Vansteelandt and Daniel (2017). The exact decomposition of the total effect into the direct effect and indirect effects via each mediator is presented, and the existing estimation method is briefly reviewed. Next, interventional effect models are introduced, and the estimation procedures, using either inverse weighting (by the counterfactual mediator density or mass functions), or Monte Carlo integration, are described. The proposed methods are then assessed via extensive simulation studies. The methods are utilized to assess the effect of self-efficacy on fatigue that is possibly mediated by different post-traumatic stress disorder symptoms, and whether the indirect effects vary across different levels of negative coping, among health care workers in China during the COVID-19 outbreak (Hou et al. 2020). The R scripts used to implement the proposed estimation procedures in carrying out the simulation studies, and the illustration, are freely available online.¹ More user-friendly functions in an R package for applied researchers are under development. A brief discussion is provided in the conclusion.

Interventional effects

Notation for counterfactual mediators and potential outcomes

Consider the setting with an exposure A , multiple mediators $M_1, \dots, M_t$, and an outcome Y . For $s = 1, \dots, t$, let $M_{s,a}$ denote the (individual) counterfactual for M_s when A is set to a . Let $\tilde{M}_{s,a|L}$ denote a random draw from the *marginal* counterfactual mediator distribution that does not depend on any other mediators (given baseline covariates L) when exposure A is set to a ; i.e., $\tilde{M}_{s,a|L} \sim F(M_{s,a}|L)$, where $F(\cdot)$ denotes a cumulative distribution function. It is assumed that all such covariates L are themselves unaffected by exposure or any mediators; otherwise they should be included among the set of mediators. For notational simplicity, write $\tilde{M}_{s,a|L}$ as $\tilde{M}_{s,a}$ where the dependence on L is implied. Hence the counterfactual mediators $\tilde{M}_{s,a} = M_{s,a}$ are equivalent when the covariates L are sufficiently rich so that the distribution $F(M_{s,a}|L)$ is degenerate with support (i.e., non-zero

¹ <https://github.com/wwloh/interventional-effects>.

probability) at only the (counterfactual) value of $M_{s,a}$; otherwise $\tilde{M}_{s,a}$ will differ from $M_{s,a}$. Let $Y_{a m_1 \dots m_t}$ denote the potential outcome for Y if, possibly counter to fact, A is set to a , and each mediator M_s is set to the value m_s , $s = 1, \dots, t$. For example, the potential outcome for Y if exposure $A = a^{(0)}$ and when $M_s = \tilde{M}_{s,a^{(s)}}$, $s = 1, \dots, t$, is thus $Y_{a^{(0)} \tilde{M}_{1,a^{(1)}} \dots \tilde{M}_{t,a^{(t)}}}$. Let $\tilde{\mathbf{M}}_a = (\tilde{M}_{1,a}, \dots, \tilde{M}_{t,a})$ denote a random draw of the counterfactual mediators from their joint distribution under exposure $A = a$ (given L); i.e., $\tilde{\mathbf{M}}_a \sim F(M_{1,a}, \dots, M_{t,a} | L)$. Here and throughout (counterfactual) mediators are written in bold if and only if referring to (values from) their joint distribution. Note that both marginal and joint (counterfactual) mediator distributions are *conditional* on the baseline covariates L , although we omit the dependence for notational convenience. Let $Y_{a^{(0)} \tilde{\mathbf{M}}_{a^{(1)}}}$ denote the potential outcome if exposure A is set to $a^{(0)}$ and when the values of the mediators are jointly drawn from $F(M_{1,a^{(1)}}, \dots, M_{t,a^{(1)}} | L)$.

Definition and decomposition of interventional effects

In this section, we define the interventional (in)direct effects and describe possible decompositions for a binary exposure A following Vansteelandt and Daniel (2017). The average potential outcomes (hereafter termed “estimands”) among individuals with a fixed value of the baseline covariates $L = l$ are defined by:

$$E(Y_{a^{(0)} \tilde{\mathbf{M}}_{a^{(1)}}} | L = l) = \int E(Y_{a^{(0)} m_1 \dots m_t} | L = l) dF(M_{1,a^{(1)}} = m_1, \dots, M_{t,a^{(1)}} = m_t | L = l),$$

and

$$E(Y_{a^{(0)} \tilde{M}_{1,a^{(1)}} \dots \tilde{M}_{t,a^{(t)}}} | L = l) = \int E(Y_{a^{(0)} m_1 \dots m_t} | L = l) dF(M_{1,a^{(1)}} = m_1 | L = l) \dots dF(M_{t,a^{(t)}} = m_t | L = l).$$

For notational simplicity, we henceforth write the conditional expectation as $E_L(\cdot) = E(\cdot | L)$. Let $g(\cdot)$ denote a user-specified canonical link function. Let Y_a denote the potential outcome for Y if, possibly counter to fact, A is set to a . Define the total effect as $g\{E_L(Y_1)\} - g\{E_L(Y_0)\}$, which can be decomposed into a direct effect and an indirect effect, respectively:

$$DE_L = g\{E_L(Y_{1\tilde{\mathbf{M}}_1})\} - g\{E_L(Y_{0\tilde{\mathbf{M}}_1})\}, \quad (1)$$

$$IE_L = g\{E_L(Y_{0\tilde{\mathbf{M}}_1})\} - g\{E_L(Y_{0\tilde{\mathbf{M}}_0})\}. \quad (2)$$

The direct effect describes the difference between the (transformed) estimands under different exposures, while holding the (joint) distribution of the (counterfactual) mediators fixed under exposure $A=1$, in the subgroup defined by the covariate(s) L . The indirect effect is the difference between the (transformed) estimands when the (counterfactual) mediators' (joint) distribution is shifted from exposure to control, among individuals whose exposure levels are set to control in the subgroup defined by L . The (joint) indirect effect (2) can be further decomposed into separate indirect effects via each mediator, and an indirect effect via the mutual dependence of the mediators. Define the indirect effect via the s -th mediator M_s , $s = 1, \dots, t$, as:

$$IE_{Ls} = g\{E_L(Y_{0\tilde{M}_{1,1} \dots \tilde{M}_{s-1,1} \tilde{M}_{s,1} \tilde{M}_{s+1,0} \dots \tilde{M}_{t,0}})\} - g\{E_L(Y_{0\tilde{M}_{1,1} \dots \tilde{M}_{s-1,1} \tilde{M}_{s,0} \tilde{M}_{s+1,0} \dots \tilde{M}_{t,0}})\}. \quad (3)$$

The exposure $a^{(s)}$ for the s -th mediator takes the value 1 in the first term, and the value 0 in the second term. This indirect effect is therefore the difference between the (transformed) estimands when the s -th mediator's (marginal) distribution is shifted from one exposure level to another, while holding the distributions of all other mediators fixed among individuals with covariate value L whose exposure levels are set to control. By shifting the marginal distribution of the counterfactual mediator in question - averaged over all other mediators - this indirect effect is readily interpreted as the combined effect of exposure on the mediator, along all underlying causal pathways (possibly via other mediators), and the effect of the mediator on the outcome. Moreno-Betancur et al. (2020) describe how the shifts in the mediators' distributions can be interpreted as hypothetical interventions in a future “target” trial. Such an interpretation avoids the notion of underlying path-specific effects among the mediators, because any future (randomized) intervention on a mediator (alone)

may be designed so that all other mediators are unaffected by that intervention. Because the (possibly arbitrary) mediator indices are used merely as labels and need not imply any assumed causal ordering, a sensitivity analysis can be readily carried out by permuting the mediator indices and assessing the indirect effects under each permutation. We conduct such a sensitivity analysis in the applied example.

The total, direct and joint indirect effects are defined using potential outcomes where the mediators are drawn from their joint (counterfactual) distribution, whereas the separate indirect effects via each mediator are defined using potential outcomes where the mediators are drawn from their respective marginal (counterfactual) distribution. This distinction leads to the following difference between the joint indirect and the sum of the separate indirect effects, where the latter equals:

$$\sum_{s=1}^t \text{IE}_{Ls} = g\{E_L(Y_{0\tilde{M}_{1,1}\dots\tilde{M}_{t,1}})\} - g\{E_L(Y_{0\tilde{M}_{1,0}\dots\tilde{M}_{t,0}})\}. \quad (4)$$

In general, the sum (4) need not equal the (joint) indirect effect (2). The difference is termed the indirect effect via the *mutual dependence of the mediators*, and is defined as:

$$[g\{E_L(Y_{0\tilde{M}_1})\} - g\{E_L(Y_{0\tilde{M}_0})\}] - [g\{E_L(Y_{0\tilde{M}_{1,1}\dots\tilde{M}_{t,1}})\} - g\{E_L(Y_{0\tilde{M}_{1,0}\dots\tilde{M}_{t,0}})\}]. \quad (5)$$

This indirect effect is an important component in the decomposition of the joint indirect effect: it describes the mediated effect of exposure on outcome when the relationships between the mediators, and their subsequent effects on the outcome, differ for different exposure levels, so that the indirect effect via the mediators cannot be considered separately through any lone mediator. For example, under linear mean models for the mediators and the outcome, this indirect effect is non-zero if (i) there is non-zero mediator-mediator interaction in the outcome model; and (ii) the covariance of the mediators differs with exposure. Closed form expressions of the interventional direct and indirect effects under (correctly) assumed linear models for the mediators and the outcome when there are two mediators are provided in Section B in the Supplementary material.

Identification of interventional effects

In general, the observed (baseline) covariates L used to define subgroups for the heterogeneous interventional effects may be confounders of the exposure-outcome, mediator(s)-outcome, and exposure-mediator(s) relations. Again, we note that all confounders in L are assumed to be unaffected by the exposure or any mediators, otherwise any exposure-induced confounders should be listed as additional mediators. Identification of the interventional effects defined above therefore requires the following assumptions (Vansteelandt and Daniel 2017):

- the effect of exposure A on outcome Y is unconfounded conditional on L , i.e.,

$$Y_{am_1\dots m_t} \perp\!\!\!\perp A|L \quad \forall a, m_1, \dots, m_t; \quad (6)$$

- the effect of all mediators $M_1, \dots, M_t$ on outcome Y is unconfounded conditional on A and L , i.e.,

$$Y_{am_1\dots m_t} \perp\!\!\!\perp (M_1, \dots, M_t)|(A = a, L) \quad \forall a, m_1, \dots, m_t; \quad (7)$$

- the effect of exposure A on all mediators is unconfounded, i.e.,

$$(M_{1,a}, \dots, M_{t,a}) \perp\!\!\!\perp A|L \quad \forall a. \quad (8)$$

Two additional assumptions ensure that the counterfactual mediators and potential outcomes can be identified from observed data. The first is that *causal consistency* holds, i.e., that the counterfactual mediators and potential outcomes equal their observed values when the hypothetical exposure and mediator values correspond to their observed values; i.e., $M_s = M_{s,A}$, $s = 1, \dots, t$, and $Y = Y_{AM_1\dots M_t}$. The second is that *positivity* holds, i.e., that each individual has (i) non-zero probabilities of being assigned to either the exposure or control group, given the covariates, e.g., $\Pr(A = a|L) > 0$, $a = 0, 1$; and (ii) non-zero mediator densities, conditional on treatment and covariates, e.g., $f(M_1 = m_1, \dots, M_t = m_t|A = a, L) > 0$ for all values of $a, m_1, \dots, m_t$.

Existing estimation methods

In this section, we briefly review existing estimation methods for interventional (in)direct effects as described in Vansteelandt and Daniel (2017). Estimation requires specifying an appropriate mean model for the outcome and a model for the mediators' joint distribution. The outcome mean model is conditional on the mediators $M_1, \dots, M_p$, exposure A and all covariates L . The model for the mediators' joint distribution is conditional on the exposure A and all covariates L ; the (implied) model for each mediator's marginal distribution can be obtained by (numerically) averaging over the distribution of the other mediators. For each specific interventional (in)direct effect in turn, the (counterfactual) mediators are randomly sampled from their distributions under the corresponding exposure levels. The potential outcomes are predicted given the sampled mediator values using the fitted model. Monte Carlo integration is carried out by making repeated stochastic draws of the mediators, then averaging over the predicted potential outcomes across all mediator draws. However, to estimate the aforementioned heterogeneous indirect effects, the potential outcomes would be averaged only among the subset of individuals who share the same value(s) of the observed covariate(s) L . Repeating this calculation for each subgroup in turn can quickly become unwieldy in practice, especially when there are more than two mediators or more than a few unique covariate values, or both. Furthermore, it is unfeasible to carry out these calculations when the covariates are continuous and only a few individuals are in a subgroup defined by the same covariate value(s). In the next section, we propose an approach to simplify estimation of heterogeneous interventional effects.

Interventional effect models

Effect models

We now describe the interventional effect models that facilitate estimating the interventional (in)direct effects proposed in this article. We adopt the same functional form as natural effect models (Lange, Vansteelandt, and Bekaert 2012; Steen et al. 2017) that generalize marginal structural models (Robins 2000) to express the mean potential outcomes in terms of hypothetical exposure levels for each mediator. The parameters in natural effect models describe differences between the estimands that exactly encode the natural (in)direct effects. But unlike natural effect models that focus on decomposing natural indirect effects, we will use interventional effect models to parametrize the aforementioned interventional (in)direct effects so as to allow simultaneous estimation.

For pedagogical purposes, we will consider the setting with two mediators M_1 and M_2 , and two covariates L_1 and L_2 , and defer models for settings with more mediators and covariates to the simulation studies and the applied example. An (interventional) effect model that encodes (in)direct effects that are modified only by the covariate L_1 is:

$$\begin{aligned} gE_L\{Y_{am_1m_2}\} = & \mu_0 + \{\theta_1 a^{(1)} + \theta_2 a^{(2)} + \theta_{1c} a^{(1)} L_1 + \theta_{2c} a^{(2)} L_1\} (1 - J) + \mu_1 L_1 + \mu_2 L_2 \\ & + \{\mu_{0J} + \gamma_0 a^{(0)} + \gamma_{0c} a^{(0)} L_1 + \gamma_1 a^{(1)} \mathbb{1}(a^{(1)} = a^{(2)}) + \gamma_{1c} a^{(1)} \mathbb{1}(a^{(1)} = a^{(2)}) L_1\} J. \end{aligned} \quad (9)$$

For notational simplicity here and throughout, we will use the subscripts on the left hand side of (9) merely to emphasize that the potential outcomes for Y , and not the observed outcomes Y , are the dependent variables in the effect models. The indicator function $\mathbb{1}(B)$ takes value 1 when event B is true, or 0 otherwise. The interactions between the mediators' exposure levels and the covariate L_1 allow for (in)direct effects to vary across different values of L_1 (The subscripts c in the coefficients θ_{1c} , θ_{2c} , γ_{0c} , γ_{1c} emphasize the interactions that encode the effect modification by the covariate of interest). The indicator J takes value 0 when the potential outcomes are defined using mediators that are randomly drawn from their respective marginal (counterfactual)

distributions, or 1 when drawn from their joint (counterfactual) distribution. When $J=0$, the indirect effect via M_1 is encoded by $(\theta_1 + \theta_{1c}L_1)$ for some fixed value of L_1 . Similarly, the indirect effect via M_2 is encoded by $(\theta_2 + \theta_{2c}L_1)$ for some fixed value of L_1 . The indirect effect measures are therefore modified by the covariate L_1 when the parameters θ_{1c} or θ_{2c} are non-zero. The sum of the indirect effects via both mediators is readily obtained by summing the parameters $\sum_{s=1}^2 (\theta_s + \theta_{sc}L_1)$. The joint indirect effect is defined using potential outcomes where the mediators are drawn from their joint (counterfactual) distribution under exposure level $A = a^{(1)}$, hence restricting the exposure levels for both mediators to be the same when $J=1$, i.e., $a^{(1)} = a^{(2)}$. The joint indirect effect is thus encoded by $(\gamma_1 + \gamma_{1c}L_1)$ for some fixed value of L_1 . The indirect effect via the mutual dependence of the mediators is the difference between the joint indirect effect and sum of the separate indirect effects, and is thus encoded by $(\gamma_1 + \gamma_{1c}L_1) - \sum_{s=1}^2 (\theta_s + \theta_{sc}L_1)$. There is no main effect for the exposure $a^{(0)}$ alone when $J=0$ because the indirect effects defined in this paper are fixed at $a^{(0)} = 0$. The direct effect is encoded by the parameters $(\gamma_0 + \gamma_{0c}L_1)$ for some fixed value of L_1 when $J=1$. Lastly, the total effect is the sum of the direct and joint indirect effects, which is simply the sum of the parameters $(\sum_{j=0}^1 \gamma_j + \gamma_{jc})$. Because interventional effect models are conditional mean models for the potential outcomes, main effects for both L_1 and L_2 are included in the effect model to adjust for confounding, even when the association between L_2 and the (potential) outcome is not of primary interest (Vansteelandt, Bekaert, and Lange 2012). Interaction terms $J \times L_1$ and $J \times L_2$ may also be included in the effect model to allow the effect of each confounder on the (potential) outcome to differ depending on whether the (counterfactual) mediators are randomly drawn from either their respective marginal distributions ($J=0$), or their joint distribution ($J=1$).

Estimation via inverse weighting

Estimators of the interventional direct and indirect effects can be obtained by fitting the effect model using weights that are (inversely) proportional to the counterfactual mediator density or mass functions. For each individual, duplicated data is constructed using different levels of the exposure $a^{(0)}$, $a^{(1)}$, $a^{(2)}$, under either marginal ($J=0$) or joint ($J=1$) distributions, for the counterfactual mediators. The posited effect model is then fitted to the duplicated data using weighted regression. The procedure is as follows:

- A1. Fit a propensity score model for exposure conditional on all observed baseline confounders to the observed data. This is required to adjust for exposure-outcome and exposure-mediator(s) confounding toward satisfying the identifying assumptions (6) and (8). For example, a logistic regression model may be:

$$\text{logit}\{\Pr(A = 1|L)\} = \beta_0 + \beta_1 L,$$

where $\text{logit}(x) = \log\{x/(1-x)\}$. Let $\hat{p} = \text{expit}(\hat{\beta}_0 + \hat{\beta}_1 L)$ denote the (individual) predicted probability of receiving the observed exposure based on the maximum likelihood estimates (MLEs) of the parameters $\hat{\beta}_0$ and $\hat{\beta}_1$, where $\text{expit}(x) = \exp(x)/[1 + \exp(x)]$. Calculate the estimated weight for each individual as $\hat{w}^a = A/\hat{p} + (1-A)/(1-\hat{p})$.

- A2. Within each observed exposure group $A = a$, fit a group-specific model for the joint density or mass function of the (counterfactual) mediators, conditional on all observed baseline confounders L , to the observed data. The confounders L are included in the mediator models to adjust for mediator(s)-outcome confounding toward satisfying the identifying assumption (7). Denote the resulting estimated density by $\hat{f}^a(\mathbf{M}|L)$, $a = 0, 1$. For example, suppose that M_1 is a binary mediator and M_2 is a normally-distributed mediator. The joint density can be factorized as $\hat{f}^a(\mathbf{M}|L) = \hat{f}^a(M_2|M_1, L)\hat{f}^a(M_1|L)$. A logistic regression model for M_1 among individuals with exposure $A = a$ may be:

$$\text{logit}\{\Pr(M_1 = 1|A = a, L)\} = \alpha_{10}^a + \alpha_{11}^a L,$$

where the superscript a in the regression coefficients denote their dependence on $A = a$. The implied estimated distribution is therefore a Bernoulli distribution with probability of success

$\hat{E}^a(M_1|L) = \text{expit}(\hat{\alpha}_{10}^a + \hat{\alpha}_{11}^a L)$, where $\hat{\alpha}_{10}^a$ and $\hat{\alpha}_{11}^a$ are the MLEs of the parameters. A linear regression model for M_2 , conditional on M_1 and L , among individuals with exposure $A = a$ may be:

$$E^a(M_2|M_1, L) = \alpha_{20}^a + \alpha_{21}^a M_1 + \alpha_{22}^a L, \quad V^a(M_2|M_1, L) = (\sigma_2^a)^2,$$

where $E^a(X|U)$ and $V^a(X|U)$ respectively denote the (conditional) expectation and variance of a random variable X given U and exposure $A = a$. The implied estimated distribution is therefore a Normal distribution with mean $\hat{\alpha}_{20}^a + \hat{\alpha}_{21}^a M_1 + \hat{\alpha}_{22}^a L$ and (constant) variance $(\hat{\sigma}_2^a)^2$, where $\hat{\alpha}_{20}^a, \hat{\alpha}_{21}^a, \hat{\alpha}_{22}^a$ and $(\hat{\sigma}_2^a)^2$ are the MLEs of the parameters.

Let $\hat{f}^a(M_s|L)$, $a = 0, 1$, denote the marginal density of each mediator M_s , $s = 1, 2$, unconditional on all other mediators but given the baseline confounders L , as implied by the mediators' joint distribution $\hat{f}^a(\mathbf{M}|L)$. Continuing the above example, it follows from the laws of total expectation and of total variance that M_2 given L is Normally distributed with mean $\hat{\alpha}_{20}^a + \hat{\alpha}_{21}^a \hat{E}^a(M_1|L) + \hat{\alpha}_{22}^a L$ and (constant) variance $(\hat{\sigma}_2^a)^2 + (\hat{\alpha}_{21}^a)^2 \hat{V}^a(M_1|L)$, where $\hat{V}^a(M_1|L) = \hat{E}^a(M_1|L)\{1 - \hat{E}^a(M_1|L)\}$ for binary M_1 .

- A3. Construct the duplicated data for each individual as shown in Table 1. In the first row, set all the hypothetical exposure levels to 0; e.g., $a^{(0)} = a^{(1)} = \dots = a^{(t)} = 0$. In each row $s = 2, \dots, t + 1$, set the hypothetical exposure levels as $a^{(k)} = 1$ if $k \in \{1, \dots, s - 1\}$, or $a^{(k)} = 0$ otherwise. The difference between the estimands in rows s and $s - 1$ therefore corresponds to the interventional indirect effect via the mediator M_{s-1} ; e.g., the indirect effect via M_2 corresponds to the difference between the estimands in the third and second rows. Similarly, the difference between the estimands in rows 1 and $t + 1$ corresponds to the sum of the interventional indirect effects via the distinct mediators. In the last two rows, set the hypothetical exposure levels $a^{(0)}$ to the observed value A ; in the penultimate row set $a^{(1)} = \dots = a^{(t)} = 1 - A$, and in the last row set $a^{(1)} = \dots = a^{(t)} = A$. There are therefore a total of $t + 3$ rows for t mediators.

- A4. For each row in the duplicated data, calculate the weight:

$$\hat{w}_i(a^{(0)}, \dots, a^{(t)}, J) = \mathbb{1}(a^{(0)} = A) \hat{w}_i^a(a^{(1)}, \dots, a^{(t)}, J),$$

$$\text{where } \hat{w}_i^m(a^{(1)}, \dots, a^{(t)}, J) = \begin{cases} \prod_{s=1}^t \hat{f}^{a^{(s)}}(M_s|L) / \hat{f}^A(\mathbf{M}|L), & J = 0; \\ \hat{f}^{a^{(1)}}(\mathbf{M}|L) / \hat{f}^A(\mathbf{M}|L), & J = 1. \end{cases}$$

- A5. Fit the effect model to the duplicated data in Table 1 with the (observed) outcomes and computed weights from the previous step using weighted regression.

Nonparametric bootstrap confidence intervals (Efron and Tibshirani 1994) may be constructed by randomly resampling observations with replacement and repeating all the steps for each bootstrap sample.

In step A2, separate exposure group-specific mediator models may potentially be simpler to specify (correctly) than models with interaction term(s) between the covariate(s) and exposure. Furthermore, it allows

Table 1: Duplicated data for each individual with five rows when there are two mediators for a binary exposure A , used to estimate the parameters in an (interventional) effect model. The (counterfactual) mediators are randomly drawn from either their respective marginal distributions ($J=0$) or their joint distribution ($J=1$). The observed covariates L (including L) are omitted for simplicity.

$a^{(0)}$	$a^{(1)}$	$a^{(2)}$	J	$\hat{w}_i^m(a^{(1)}, \dots, a^{(t)}, J)$	Y
0	0	0	0	$\hat{f}^0(M_1 L)\hat{f}^0(M_2 L)/\hat{f}^A(\mathbf{M} L)$	Y
0	1	0	0	$\hat{f}^1(M_1 L)\hat{f}^0(M_2 L)/\hat{f}^A(\mathbf{M} L)$	Y
0	1	1	0	$\hat{f}^1(M_1 L)\hat{f}^1(M_2 L)/\hat{f}^A(\mathbf{M} L)$	Y
A	$1 - A$	$1 - A$	1	$\hat{f}^{1-A}(\mathbf{M} L)/\hat{f}^A(\mathbf{M} L)$	Y
A	A	A	1	1	Y

for exposure effects on the (joint) mediators' distributions (including their residual variances), and not merely their means (as implied by interaction terms in a single mediator model for both exposure groups). We emphasize that the factorization of the mediators' joint density is not predicated on any (correctly) assumed causal ordering among the mediators; e.g., $\hat{f}^a(\mathbf{M}|L) = \hat{f}^a(M_1|M_2, L)\hat{f}^a(M_2|L)$ may be merely used instead. The proposed estimation procedure follows Lange, Rasmussen, and Thygesen (2013), by fitting the effect model to the observed outcomes using weighted regression techniques. Consistent estimation is similarly predicated on correctly specifying (i) a propensity score model that converges in probability to its true value, and (ii) models for the (joint) distribution of the mediators, conditional on exposure and covariates, that are unbiased for the counterfactual distribution, i.e., $f^a(\mathbf{M}|L)$ converges in probability to its counterfactual distribution $f(M_a|L)$, for $a = 0, 1$. A proof, which includes the motivation for the form of the weights, is provided in Section A in Supplementary material. However, inverse weighting by the joint mediator density or mass functions (Hong 2010; Lange, Vansteelandt, and Bekaert 2012; Steen et al. 2017) may result in highly variable or volatile weights that induce estimators with non-negligible finite sample bias under certain situations. In particular, when the mediators are (strongly) associated with one another, given A and L , due to either causal effects on each other or unobserved confounding, or both, estimators of the separate indirect effects through the mediators may be biased. We demonstrate the biases empirically in a simulation study.

Estimation via Monte Carlo integration

To avoid potential biases resulting from highly variable or volatile weights, we propose a second estimator that does not require inverse weighting. Instead, Monte Carlo draws of the counterfactual mediators are made at different exposure levels in a duplicated dataset similar to Table 1. A specified outcome model is used to predict potential outcomes given the counterfactual mediators and observed confounders in the duplicated data. Estimators of the interventional (in)direct effects are then obtained by fitting the posited effect model to the duplicated data, using either ordinary least squares or maximum likelihood estimation. The procedure is as follows:

- B1. Within each observed exposure group $A = a$, fit a group-specific outcome model conditional on all mediators and observed baseline confounders, e.g., $E(Y|A = a, M_1, M_2, L)$, to the observed data. Each separate outcome model can be expressed as a function of its inputs, e.g., $E(Y|A = a, M_1 = m_1, M_2 = m_2, L) = h^a(m_1, m_2, L)$, where $h^a(\cdot)$ is a user-specified function with the superscript a denoting its dependence on $A = a$. Denote the estimated functions by $\hat{h}^a(m_1, m_2, L)$, $a = 0, 1$. For example, a logistic regression model for a binary outcome among individuals with exposure $A = a$ is:

$$\text{logit}\{\Pr(Y = 1|A = a, M_1, M_2, L)\} = \beta_0^a + \beta_1^a M_1 + \beta_2^a M_2 + \beta_{11}^a M_1 L + \beta_{21}^a M_2 L + \beta_1^a L.$$

Table 2: Duplicated data for each individual with five rows when there are two mediators for a binary exposure A , used to estimate the parameters in an (interventional) effect model. The (counterfactual) mediators are randomly drawn from either their respective marginal distributions ($J=0$) or their joint distribution ($J=1$). The asterisk denotes a stochastically imputed counterfactual mediator. The observed covariates L (including L) are omitted for simplicity.

$a^{(0)}$	$a^{(1)}$	$a^{(2)}$	J	$\tilde{M}_{1, a^{(1)}}$	$\tilde{M}_{2, a^{(2)}}$	$h^{a^{(0)}}(m_1 = M_{1, a^{(1)}}, m_2 = M_{2, a^{(2)}}, L)$
0	0	0	0	$M_{1,0}^*$	$M_{2,0}^*$	$\hat{h}^0(m_1 = M_{1,0}^*, m_2 = M_{2,0}^*, L)$
0	1	0	0	$M_{1,1}^*$	$M_{2,0}^*$	$\hat{h}^0(m_1 = M_{1,1}^*, m_2 = M_{2,0}^*, L)$
0	1	1	0	$M_{1,1}^*$	$M_{2,1}^*$	$\hat{h}^0(m_1 = M_{1,1}^*, m_2 = M_{2,1}^*, L)$
$1 - A$	A	A	1	$\mathbf{M} = (M_1, M_2)$		$\hat{h}^{1-A}(m_1 = M_1, m_2 = M_2, L)$
A	A	A	1	$\mathbf{M} = (M_1, M_2)$		Y

- B2. Following step A2, within each observed exposure group $A = a$, fit a group-specific model for the marginal density of each mediator M_s , $s = 1, 2$, conditional on all baseline confounders L , to the observed data. Denote the resulting estimated distribution by $\hat{F}^a(M_s|L)$.
- B3. Construct the duplicated data for each individual as shown in Table 2, which differs from Table 1 only in the last two rows. In the last two rows of Table 2, set the hypothetical exposure levels $a^{(1)} = \dots = a^{(t)}$ to the observed value A ; in the penultimate row set $a^{(0)} = 1 - A$, and in the last row set $a^{(0)} = A$. There are therefore a total of $t + 3$ rows for t mediators.
- B4. For each row $s = 1, \dots, t + 1$ where $J=0$, randomly sample the counterfactual mediator values (denoted by asterisks in the superscripts) from their respective (estimated) marginal distributions $\hat{F}^{a^{(s)}}(M_s|L)$.
- B5. Using the separate fitted outcome models for $A=1$ and $A=0$, impute the expected potential outcomes in each row as predictions $\hat{h}^{a^{(0)}}(m_1 = \tilde{M}_{1,a^{(1)}}, m_2 = \tilde{M}_{2,a^{(2)}}, L)$, depending on whether $a^{(0)} = 0$ or $a^{(0)} = 1$. In the last row where $a^{(0)} = A$ and $J=1$, set the potential outcome to its observed value.
- B6. Fit the effect model to the duplicated data for the observed sample.

Nonparametric bootstrap confidence intervals may be constructed by randomly resampling observations with replacement and repeating all the steps for each bootstrap sample.

No (propensity score) model for the exposure (given the observed confounders L) is required, even when exposure is non-randomly assigned, because including the observed confounders L in the outcome and mediator models suffices to adjust for exposure-outcome and exposure-mediator(s) confounding toward satisfying the identifying assumptions (6) and (8). In step B1, separate exposure group-specific outcome mean models are recommended for the same reason as separate exposure group-specific mediator models in step A2. The models can potentially be simpler to specify (correctly) than a single model that includes covariate-exposure and (possibly numerous) exposure-mediator(s) interaction effect(s) on the observed outcome. Assuming that the residuals from the mediator models fitted to the observed data in step B2 are independent of all variables, an alternative to random draws of the mediators from specified parametric distributions in step B4 is to resample (with replacement) the residuals, then attach them to the predicted (counterfactual) mediators. In the last two rows of Table 2, the observed values of the mediators are viewed as being drawn from their (possibly unknown) joint distribution under the observed exposure level. Using the observed values has the benefit of being computationally efficient (by avoiding sampling the counterfactual mediators). Furthermore, the risks of biases due to model misspecification can be avoided, because assuming a model for the joint distribution of the (counterfactual) mediators, conditional on the confounders and (hypothetical) exposure, is unnecessary for the proposed Monte Carlo estimator. Because our interest is only in the expected potential outcome in each row, steps B4 and B5 may be repeated e.g., 100 times, to average over the (counterfactual) distributions of the mediators using Monte Carlo integration. In each row of the duplicated data, the averaged (predicted) potential outcomes over all Monte Carlo draws can then be employed as the imputed potential outcome for that individual.

The proposed imputation-based procedure follows the estimation strategy of Vansteelandt, Bekaert, and Lange (2012) for a single mediator. It is similar to the imputation-based strategies for G-computation (Snowden, Rose, and Mortimer 2011), and shares their virtues and limitations (Vansteelandt and Keiding 2011). Consistent estimation is predicated on correctly specifying (i) an outcome model conditional on exposure, mediators, and covariates that is unbiased for the marginal structural mean model, i.e., $h^a(m_1, \dots, m_t, L)$ converges in probability to $E(Y_{am_1 \dots m_t}|L)$; and (ii) models for the (marginal) distributions of the mediators, conditional on exposure and covariates, that are unbiased for their counterfactual distributions, i.e., for $s = 1, \dots, t$, $f^a(M_s|L)$ converges in probability to its counterfactual distributions $f(M_{s,a}|L)$, for $a = 0, 1$. A proof is provided in Section A in Supplementary material. However, its simplicity may belie the same difficulty facing multiple imputation estimators for missing data analyses: specifying a model for the outcome that is “congenial” (Meng 1994), or “coherent,” with the effect model of interest. An assumed (parametric) outcome model that is uncongenial with the posited effect model may result in extrapolation (or “misspecification”) bias, although uncongeniality can

be avoided in principle by using saturated or rich outcome models (Vansteelandt, Bekaert, and Lange 2012). For example, when the effect model is linear, i.e., $g(\cdot)$ being the identity link, an outcome model that reflects the structure of the effect model may be obtained by (i) considering only the submodel where $J=0$, (ii) replacing $a^{(s)}$ with M_s , $s = 1, \dots, t$, then (iii) adding mediator-mediator(-covariate) interaction terms, such as $M_1 \times M_2$ and $M_1 \times M_2 \times L$, as well as an exposure A and covariate-exposure interaction(s) $A \times L$. We provide such an example in the illustration. But in practice, it may be unfeasible to fit to the observed data saturated models when there are continuous covariates or mediators, or models with all possible interaction and higher-order terms when there are high-dimensional (continuous) covariates or mediators. Furthermore, specifying an outcome model with the correct structure as the posited effect model can be difficult when the latter is nonlinear due to noncollapsibility; see e.g., VanderWeele and Vansteelandt (2010) and Tchetgen Tchetgen (2014) when using logistic regression for binary outcomes. We evaluate the impact of misspecifying the outcome model on the effect estimates empirically in a simulation study, and defer formal assessment of the sensitivity of the estimators to model misspecification to future work.

Simulation studies

Three simulation studies across 45 different settings were conducted to empirically assess the operating characteristics of the proposed estimators using either inverse weighting or Monte Carlo integration. Details of the data-generating procedures are deferred to Section C in Supplementary material. To provide an overview, in study 1, we considered a simple setting with two mediators and a single confounder, with no interaction terms in the data-generating models, so that the (in)direct effects were the same across different levels of the confounder. The objective of this study was to compare the empirical biases of the two estimation procedures, especially the impact of the (in)stability of the weights for the inverse weighting estimator on any potential biases. To demonstrate that the mediators' joint density used in the inverse weighting estimator did not require correctly ordering the mediators (e.g., $M_1 \rightarrow M_2$), we simply factorized the joint density into its conditional distributions using an incorrect causal ordering instead (e.g., $M_2 \rightarrow M_1$). In study 2, we built on study 1 to introduce an additional confounder, and an exposure-confounder interaction effect on a mediator, in the data-generating model to render separate indirect effects via each mediator that differed across different levels of that confounder. In addition to the empirical biases of the two estimation procedures, we compared the empirical coverage of their corresponding nonparametric bootstrap percentile 95% confidence intervals. In study 3, we built on both studies 1 and 2 by further including exposure-mediator and mediator-mediator interaction effects in the data-generating model. This resulted in all the indirect effects, including the effect via the mediators' mutual dependence, being modified by a confounder. In addition to assessing the empirical biases of both the inverse weighting and Monte Carlo estimators under correctly specified models, we considered a third Monte Carlo estimator that misspecified the outcome model. In all three studies we considered settings with two (normally-distributed) mediators, where one may have been affected by the other, and with an outcome that was binary in studies 1 and 3, and normally-distributed in study 2.

We summarize the results of each simulation study comparing the two proposed estimation procedures, and refer readers to the detailed findings in Section C in Supplementary material. In study 1, the inverse weighting estimators of the direct and joint indirect effects were empirically unbiased under all the considered settings, even when an incorrect ordering for factorizing the joint mediator density was chosen. The estimators of the separate indirect effects via each mediator were empirically unbiased, but only when either (i) there were no or weak causal effects among the mediators, (ii) the effect of exposure transmitted (indirectly) to the mediator in question via any other mediators was similar to the effect transmitted (directly) to the mediator. The non-negligible biases, which persisted even at a large sample size, were due to less stable weights when one of the mediators was strongly affected (indirectly only) by exposure. In contrast, the Monte Carlo estimators were empirically unbiased under all the considered settings. Note that the Monte Carlo estimators as proposed in this paper requires no assumptions about the causal ordering of the mediators, because models for only the marginal distribution of each mediator (unconditional on all other mediators), and not the mediators'

joint distribution, are required. In study 2, similar results for the inverse weighting point estimates were arrived at. Furthermore, the instability in the weights led to (possibly severe) under-coverage of the confidence intervals (far) below their nominal levels. In contrast, estimators using the Monte Carlo procedure were empirically unbiased, and the confidence intervals approximately attained the nominal coverage, under all considered settings. In study 3, similar results for the inverse weighting point estimates were again arrived at. However, when the fitted outcome model incorrectly included only main effects for the mediators, the Monte Carlo estimators were biased even at a large sample size. When the outcome model (correctly) included the mediator-mediator interaction terms, the Monte Carlo estimators displayed finite sample biases and larger variability empirically that were reduced as the sample size increased.

In conclusion, the inverse weighting estimators were susceptible to empirically biased point estimates and confidence intervals, due to less stable weights when one of the mediators was strongly affected (indirectly only) by exposure, even at a large sample size. In contrast, the Monte Carlo estimators were empirically consistent for their true effects when the (parametric) outcome model was correctly specified.

Application

Data and variables

The proposed estimation procedures are illustrated by reanalyzing a publicly available data set² from an observational study exploring the association between self-efficacy and fatigue among $n=527$ health care workers during the COVID-19 outbreak (Hou et al. 2020). The data was collected between March 13 and 20, 2020, from a cross-sectional survey of workers in Anqing City, Anhui Province, China, which borders Hubei province, the epicenter of the COVID-19 outbreak. Frontline health care workers experience the recurring strains and grief of treating patients during the pandemic which severely affects their mental health (Noguchi 2020). Hou et al. investigated whether the effect of self-efficacy on fatigue among these workers could be mediated by their self-reported post-traumatic stress disorder (PTSD) symptoms. We describe the variables merely for illustrating the proposed interventional (in)direct effects and estimation procedures in this article. Readers interested in the study details and substantive justification of the self-reported PTSD symptoms as a (set of) mediator(s) lying on the causal path between self-efficacy and fatigue are referred to Hou et al. (2020); self-efficacy was measured by 10 items each scored on a 4-point Likert scale from 1 (not true at all) to 4 (exactly true). Fatigue (outcome Y) was measured by 14 dichotomous items, each indicating the presence of a symptom of physical or mental fatigue (1 if having the symptom, 0 if no symptom reported). The PTSD symptoms were measured by 17 items, each rated on a 5-point Likert scale from 1 (not severe at all) to 5 (extremely severe). The items for the PTSD symptoms could be classified into three subscales (re-experiencing, avoidance and hyperarousal), although Hou et al. defined a single mediator using the sum scores for all 17 items. Hou et al. further assessed whether indirect effects varied across different levels of negative coping, which was measured by eight items each rated on a 4-point Likert scale from 0 (never used) to 3 (often used).

For the sole purposes of illustration, we dichotomized the sum score for self-efficacy at the empirical median by defining $A=1$ if self-efficacy was above the sample median, or 0 otherwise. Following the recommendation by Hou et al., we dichotomized the sum score for fatigue so that individuals were deemed to be fatigued ($Y=1$) if they reported having at least seven symptoms, or not fatigued ($Y=0$) if they reported fewer than seven symptoms. In this paper, we considered each PTSD subscale as a separate mediator that was measured by their respective sum of the manifest items (i.e., “sum score”). The three (continuous) mediators were therefore re-experiencing (M_1 ; five items), avoidance (M_2 ; seven items) and hyperarousal (M_3 ; five items). We defined the sum score of the eight items measuring negative coping as the (baseline) covariate L_1 that potentially modified the (in)direct effects, with higher scores indicating greater tendency to use negative

² <https://www.openicpsr.org/openicpsr/project/119156/version/V1/view>.

Figure 1: Posited causal diagram for the variables in the applied example. For simplicity, the baseline covariates L are represented by a single node, but the effects emanating from each separate variable in L were allowed to have different strengths. Under the identification assumptions (6)–(8), the covariates in L were sufficient to adjust for all confounding between A , M_s , and Y for $s = 1, 2, 3$. Notwithstanding the absence of edges among the mediators (M_1 , M_2 , M_3), the mediators may be associated due to an unknown causal ordering, or unobserved common causes, or both. Edges emanating from L are drawn as gray broken lines for visual clarity.

coping. To ease interpretation of the indirect effects, we standardized the sum scores so that the sample mean and standard deviation of negative coping were zero and one respectively. The remaining observed covariates were age, gender, marital status, education level, working experience (in years) and job rank seniority; we denoted these variables, as well as negative coping, jointly by L . We assumed that the observed covariates L were unaffected by exposure, mediators or outcome, and that they were sufficient for the identification assumptions (6)–(8) to hold. A causal diagram (Pearl 2009) representing the causal relations between the variables, and the identification assumptions, is shown in Figure 1. We emphasize that the mediator indices are arbitrary, and do not reflect any a priori causal ordering among the mediators.

Analysis without effect modification

We first posited the following effect model for the interventional direct and indirect effects without any effect modification by the covariates.

$$\begin{aligned} \text{logit}\{E_L(Y_{am_1m_2m_3})\} = & \mu_0 + \left(\sum_{s=1}^3 \theta_s a^{(s)} \right) (1 - J) + \{ \mu_{0J} + \gamma_0 a^{(0)} + \gamma_1 a^{(1)} \mathbb{1}(a^{(1)} = a^{(2)} = a^{(3)}) \} J \\ & + \mu_{age} \text{age} + \mu_{gen} \text{gender} + \mu_{mar} \text{marital status} + \mu_{edu} \text{education level} \\ & + \mu_{wor} \text{working experience} + \mu_{sen} \text{rank seniority} + \mu_{neg} \text{negative coping}. \end{aligned} \quad (10)$$

The interventional indirect effects via each mediator M_s , $s = 1, 2, 3$, were encoded by θ_s . The joint indirect effect was encoded by γ_1 , and the indirect effect via the mediators' mutual dependence was encoded by $\gamma_1 - \sum_{s=1}^3 \theta_s$. The direct effect was encoded by γ_0 . Because the exposure was not randomly assigned, we assumed the following propensity score model with main effects for all covariates in L for the inverse weighting procedure.

$$\begin{aligned} \text{logit}\{\Pr(A = 1|L)\} = & v_0 + v_{age} \text{age} + v_{gen} \text{gender} + v_{mar} \text{marital status} + v_{edu} \text{education level} \\ & + v_{wor} \text{working experience} + v_{job} \text{rank seniority} + v_{neg} \text{negative coping}. \end{aligned}$$

The joint mediator distribution was factorized as $f^a(M|L) = f^a(M_1|L)f^a(M_2|M_1, L)f^a(M_3|M_1, M_2, L)$. Again we reiterate that the factorization does not assume or imply any causal ordering among the mediators. We assumed all three mediators to be normally distributed, and fitted to the observed data the following linear regression models for the mediators within each exposure group $A = a$:

$$E(M_s|A = a, L) = \alpha_{s0}^a + \alpha_{s,age}^a \text{age} + \alpha_{s,gen}^a \text{gender} + \alpha_{s,mar}^a \text{marital status} + \alpha_{s,edu}^a \text{education level} \\ + \alpha_{s,wor}^a \text{working experience} + \alpha_{s,job}^a \text{rank seniority} + \alpha_{s,neg}^a \text{negative coping}, \quad s = 1;$$

$$E(M_s|A = a, M_1, \dots, M_{s-1}, L) \\ = \eta_{s0}^a + \sum_{k=1}^{s-1} \eta_{s,k}^a M_k + \sum_{k,l=1:k>l}^{s-1} \eta_{s,k,l}^a M_k M_l \\ + \eta_{s,age}^a \text{age} + \eta_{s,gen}^a \text{gender} + \eta_{s,mar}^a \text{marital status} + \eta_{s,edu}^a \text{education level} \\ + \eta_{s,wor}^a \text{working experience} + \eta_{s,job}^a \text{rank seniority} + \eta_{s,neg}^a \text{negative coping}, \quad s = 2, 3.$$

Constant error variances were assumed in each mediator model. The Monte Carlo estimation procedure required further assuming the following logistic regression outcome model that was fitted within each exposure group $A = a$:

$$\text{logit}\{\Pr(Y = 1|A = a, M_1, M_2, M_3, L)\} \\ = \beta_0^a + \sum_{s=1}^3 \beta_s^a M_s + \sum_{s,k=1:s>k}^3 \beta_{s,k}^a M_s M_k \\ + \beta_{age}^a \text{age} + \beta_{gen}^a \text{gender} + \beta_{mar}^a \text{marital status} + \beta_{edu}^a \text{education level} \\ + \beta_{wor}^a \text{working experience} + \beta_{job}^a \text{rank seniority} + \beta_{neg}^a \text{negative coping}.$$

Mediator-mediator interaction terms were included in the mediator and outcome models to allow the effect of each mediator to differ based on the levels of another mediator. We carried out both inverse weighting and Monte Carlo estimation procedures. The resulting weights using the former procedure were relatively unstable as shown in Figure 2. Because there were three mediators in the applied example, the duplicated data for each individual corresponding to Table 1 consisted of six rows (instead of just five when there were only two mediators). The first four rows corresponded to the (counterfactual) mediators under their respective marginal distributions, whereas the last two rows corresponded to the (counterfactual) mediators under their joint distribution.

Nonparametric percentile bootstrap 95% confidence intervals were constructed using 10,000 bootstrap samples that randomly resampled n observations with replacement and repeated the estimation procedures for each bootstrap sample. Because the estimated interventional (in)direct effects using the inverse weighting procedure differed from those using the Monte Carlo procedure, and the weights appeared to be highly variable, we only presented the 95% confidence intervals (CIs) for the Monte Carlo procedure in Table 3. The estimated interventional indirect effects of self-efficacy on fatigue that were mediated through each of the subscales for the self-reported PTSD symptoms ($\theta_s, s = 1, 2, 3$) suggested that shifting the (counterfactual) distribution of each subscale from those with self-efficacy levels above the empirical median to those with levels below the median resulted in lower fatigue on average, while holding the distributions of the other

Figure 2: Weights used in the inverse weighting procedure to estimate the interventional (in)direct effects in the applied example. No mediator-covariate interaction terms were assumed in the mediator and outcome models. The row numbers index the duplicated data corresponding to Table 1 (but with an additional row because there were three mediators) used to fit the interventional effect model. The weights were log-transformed (base 10), and only the non-zero weights for each row were then plotted.

Table 3: Estimates and 95% bootstrap (percentile) confidence intervals (“95% CI”) for the interventional (in)direct effects, as encoded by the parameters in (10), for the applied example. The estimates were obtained using either the inverse weighting (“IW”) procedure, or the imputation-based Monte Carlo (“MC”) procedure. The CIs were calculated for the MC estimator only. No mediator-covariate interaction terms were assumed in the mediator and outcome models. All results were rounded to two decimal places.

Effect	Parameter	IW	MC	95% CI for MC
Indirect via re-experiencing	θ_1	0.05	−0.01	(−0.21, 0.17)
Indirect via avoidance	θ_2	−0.53	−0.22	(−0.50, 0.07)
Indirect via hyperarousal	θ_3	0.00	−0.30	(−0.54, −0.03)
Indirect via mutual dependence of all mediators	$\gamma_1 - \sum_{s=1}^3 \theta_s$	−0.31	−0.04	(−0.25, 0.14)
Direct	γ_0	−0.56	−0.76	(−1.17, −0.44)

subscales fixed and setting self-efficacy to below the empirical median. In particular, the effect of self-efficacy on fatigue that was mediated through the hyperarousal subscale (θ_3) was estimated to be −0.30, with a 95% CI of (−0.54, −0.03). The interventional direct effect of self-efficacy on fatigue that was unmediated through the measured PTSD symptoms (γ_0) had the largest magnitude among the estimated (in)direct effects, with an estimated value of −0.76 (95% CI = (−1.17, −0.44)).

The interventional indirect effects via each mediator as defined in this article can differ depending on the (arbitrary) indices used merely to label the mediators, because the other mediators are fixed at different hypothetical exposure levels. For example, when re-experiencing is indexed as M_1 , the indirect effect via re-experiencing compares the (potential outcomes in the) first two rows of Table 1, 2. But had re-experiencing been merely indexed as M_2 instead, the indirect effect would then correspond to the second and third rows. Nonetheless, we reiterate that the conceptual interpretations of the interventional indirect effect via each mediator are invariant to the choice of indices. We carried out a sensitivity analysis, by considering each of the $3! = 6$ possible permutations of the three mediators in turn, and calculated the indirect effects (and the 95% CIs) for each permutation using the Monte Carlo procedure. Each permutation of the indices implies a different decomposition of the joint indirect effect, and subsequently, may result in different estimates of the separate indirect effects. In practice, the substantive research question may be used to determine the most relevant decomposition (of the joint indirect effect) from a policy perspective, which subsequently implies the ordering of the mediator indices (Moreno-Betancur et al. 2020). The minimum and maximum estimates (and bounds of the 95% CIs) across all the permutations are shown in Table 4. The conclusions for the indirect effects were generally unchanged across the different decompositions resulting from different permutations of the mediator indices. Among the three PTSD subscales, the estimated interventional indirect effect was strongest for hyperarousal.

Table 4: Interventional indirect effect estimates and 95% bootstrap (percentile) confidence intervals (“95% CI”) using the Monte Carlo procedure for the applied example. The minimum (“Min.”) and maximum (“Max.”) estimates, and 95% CI lower and upper bounds, across all $3! = 6$ possible permutations of the mediator indices are presented. No mediator-covariate interaction terms were assumed in the mediator and outcome models. All results were rounded to two decimal places.

Interventional indirect effect	Estimate		95% CI (lower)		95% CI (upper)	
	Min.	Max.	Min.	Max.	Min.	Max.
Re-experiencing	−0.10	0.01	−0.32	−0.21	0.15	0.24
Avoidance	−0.22	−0.12	−0.50	−0.46	0.07	0.19
Hyperarousal	−0.32	−0.30	−0.56	−0.54	−0.07	0.00

Analysis with effect modification

We further assessed whether the interventional indirect effects were modified by the baseline covariate negative coping, by positing the following effect model:

$$\begin{aligned} \text{logit}\{E_L(Y_{am_1m_2m_3})\} = & \mu_0 + \left\{ \sum_{s=1}^3 \theta_s a^{(s)} + \theta_{sc} a^{(s)} \text{negative coping} \right\} (1-J) \\ & + \{ \mu_{0J} + \gamma_0 a^{(0)} + \gamma_{0c} a^{(0)} \text{negative coping} \\ & + \gamma_1 a^{(1)} \mathbb{I}(a^{(1)} = a^{(2)} = a^{(3)}) + \gamma_{1c} a^{(1)} \mathbb{I}(a^{(1)} = a^{(2)} = a^{(3)}) \text{negative coping} \} J \\ & + \mu_{age} \text{age} + \mu_{gen} \text{gender} + \mu_{mar} \text{marital status} + \mu_{edu} \text{education level} \\ & + \mu_{wor} \text{working experience} + \mu_{sen} \text{rank seniority} + \mu_{neg} \text{negative coping}. \end{aligned} \quad (11)$$

The interventional indirect effects via each mediator M_s , $s = 1, 2, 3$ (at the sample average level of negative coping) were encoded by θ_s , and the effects were modified by negative coping when the parameters θ_{sc} were non-zero. Similarly, the joint indirect and direct effects were encoded by γ_1 and γ_0 respectively; the effect measures were modified by negative coping when γ_{1c} and γ_{0c} were non-zero respectively.

We carried out both inverse weighting and Monte Carlo estimation procedures. The inverse weighting estimators were calculated using the same fitted propensity score model as before. The linear mediator mean models, and logistic regression outcome model, were augmented with interaction terms between the mediator(s) and negative coping to allow the effects of the mediators on each other, and on the outcome, to differ based on negative coping; i.e.,

$$\begin{aligned} & E(M_s | A = a, M_1, \dots, M_{s-1}, L) \\ & = \eta_{s0}^a + \sum_{k=1}^{s-1} (\eta_{s,k}^a + \eta_{s,k,neg}^a \text{negative coping}) M_k + \sum_{k,l=1:k>l}^{s-1} (\eta_{s,k,l}^a + \eta_{s,k,l,neg}^a \text{negative coping}) M_k M_l \\ & \quad + \eta_{s,age}^a \text{age} + \eta_{s,gen}^a \text{gender} + \eta_{s,mar}^a \text{marital status} + \eta_{s,edu}^a \text{education level} \\ & \quad + \eta_{s,wor}^a \text{working experience} + \eta_{s,job}^a \text{rank seniority} + \eta_{s,neg}^a \text{negative coping}, \quad s = 2, 3; \end{aligned}$$

Table 5: Estimates and 95% bootstrap (percentile) confidence intervals (“95% CI”) for the interventional (in)direct effects, as encoded by the parameters in (11), for the applied example. The estimates were obtained using either the inverse weighting (“IW”) procedure, or the imputation-based Monte Carlo (“MC”) procedure. The CIs were calculated for the MC estimator only. All results were rounded to two decimal places.

Effect	Parameter	IW	MC	95% CI	
Indirect via re-experiencing	θ_1	0.01	0.02	−0.22	0.18
Indirect via re-experiencing (with negative coping)	θ_{1c}	−0.11	0.14	−0.20	0.33
Indirect via avoidance	θ_2	−0.55	−0.23	−0.51	0.11
Indirect via avoidance (with negative coping)	θ_{2c}	−0.22	−0.01	−0.40	0.34
Indirect via hyperarousal	θ_3	0.00	−0.33	−0.56	−0.02
Indirect via hyperarousal (with negative coping)	θ_{3c}	−0.00	−0.19	−0.42	0.15
Indirect via mutual dependence	$\gamma_1 - \sum_{s=1}^3 \theta_s$	−0.16	−0.00	−0.27	0.16
Indirect via mutual dependence (with negative coping)	$\gamma_{1c} - \sum_{s=1}^3 \theta_{sc}$	−0.11	−0.08	−0.21	0.21
Direct	γ_0	−0.63	−0.76	−1.19	−0.44
Direct (with negative coping)	γ_{0c}	0.50	0.07	−0.33	0.42

Table 6: Interventional indirect effect estimates and 95% bootstrap (percentile) confidence intervals (“CI”) using the Monte Carlo procedure for the applied example. The minimum (“Min.”) and maximum (“Max.”) estimates, and 95% CI lower and upper bounds, across all $3! = 6$ possible permutations of the mediator indices are presented. Indirect effects that were possibly modified by negative coping, and thus corresponded to the interaction terms in the effect model, were denoted using a colon (“:”). The remaining indirect effects corresponded to the main effects in the effect model. All results were rounded to two decimal places.

Interventional indirect effect	Estimate		95% CI (lower)		95% CI (upper)	
	Min.	Max.	Min.	Max.	Min.	Max.
Re-experiencing	−0.09	0.02	−0.32	−0.22	0.16	0.23
Re-experiencing: Negative coping	−0.02	0.14	−0.25	−0.20	0.25	0.33
Avoidance	−0.23	−0.13	−0.51	−0.47	0.11	0.19
Avoidance: Negative coping	−0.01	0.11	−0.42	−0.34	0.33	0.40
Hyperarousal	−0.33	−0.30	−0.57	−0.55	−0.04	0.04
Hyperarousal: Negative coping	−0.24	−0.14	−0.49	−0.42	0.15	0.25

$$\begin{aligned}
 & \text{logit}\{\Pr(Y = 1|A = a, M_1, M_2, M_3, L)\} \\
 &= \beta_0^a + \sum_{s=1}^3 (\beta_s^a + \beta_{s, \text{neg}}^a \text{negative coping}) M_s + \sum_{s,k=1:s>k}^3 (\beta_{s,k}^a + \beta_{s,k, \text{neg}}^a \text{negative coping}) M_s M_k \\
 &+ \beta_{\text{age}}^a \text{age} + \beta_{\text{gen}}^a \text{gender} + \beta_{\text{mar}}^a \text{marital status} + \beta_{\text{edu}}^a \text{education level} \\
 &+ \beta_{\text{wor}}^a \text{working experience} + \beta_{\text{job}}^a \text{rank seniority} + \beta_{\text{neg}}^a \text{negative coping}.
 \end{aligned}$$

The results from fitting the models to the observed data are displayed in Section D in Supplementary material. Because the estimated weights for the inverse weighting estimator were similarly unstable as those shown in Figure 2, we again only calculated the 95% CIs for the Monte Carlo estimator. The results displayed in Table 5 suggested that negative coping primarily modified the indirect effect via the hyperarousal subscale for the self-reported PTSD symptoms (θ_{3c}), with the estimate of θ_{3c} having the largest (absolute) magnitude among the three subscales θ_{sc} , $s = 1, 2, 3$. The estimated reduction in fatigue - due to the indirect effect of self-efficacy via hyperarousal - tended to be stronger for those with higher negative coping levels; e.g., at one standard deviation above the sample mean of negative coping, there was an estimated further reduction of -0.19 (95% CI = $(-0.42, 0.15)$). Conversely, the effect of self-efficacy that was mediated through the re-experiencing subscale was higher fatigue among those with higher negative coping levels; e.g., the estimated increase in fatigue at one standard deviation above the sample mean of negative coping was 0.14 (95% CI = $(-0.20, 0.33)$). Similar conclusions from the previous analysis were arrived at. The estimated interventional indirect effects of self-efficacy on fatigue were primarily mediated through the hyperarousal subscale. The interventional direct effect that was unmediated through the measured PTSD symptoms (γ_0) had the largest magnitude among the estimated (in)direct effects, with an estimated value of -0.76 (95% CI = $(-1.19, -0.44)$) at the sample mean of negative coping. The conclusions for the indirect effects were generally unchanged after considering the different decompositions in a sensitivity analysis, as shown in Table 6.

Discussion

In this article we introduced interventional effect models for estimating interventional indirect effects for multiple mediators that may be modified by observed baseline covariates. We described two estimation procedures for the causal parameters that index the effect models: inverse weighting by the counterfactual mediator density or mass functions, or Monte Carlo-based integration. The former approach may be better suited when the mediator weights are relatively stable, such as when the mediators are weakly (or un)affected

by exposure or each other, whereas the latter approach may be better suited when it is reasonable to assume that the parametric outcome model is correct.

There are several avenues of possible future research related to mediation analyses with multiple mediators using interventional effects developed in this paper. While the interventional indirect effects defined in this paper have been fixed at the hypothetical exposure level $a^{(0)} = 0$, in principle, the indirect effects can be defined using $a^{(0)} = 1$. More general definitions of the interventional indirect effects permit exploring different decompositions, by including interactions between the hypothetical exposures, e.g., $a^{(0)}$ with either $a^{(1)}$ or $a^{(2)}$, or both, in the posited effect model. Similar to the natural effect models of Lange, Vansteelandt, and Bekaert (2012) and Steen et al. (2017), the interactions encode (in)direct effects that may be modified by the values of the mediators (and exposure) under different (hypothetical) exposure levels. For example, augmenting the effect model in (10) with an additional term $\gamma_{01}a^{(0)}a^{(1)}\mathbb{I}(a^{(1)} = a^{(2)} = a^{(3)})J$ permits investigating two different decompositions of the total effect (concurrently within the same effect model). Under the first decomposition, the direct and (joint) indirect effects, as previously defined in (1) and (2), are respectively encoded by: $\gamma_0 + \gamma_{01} = g\{E_L(Y_{1\bar{M}_1})\} - g\{E_L(Y_{0\bar{M}_1})\}$, $\gamma_1 = g\{E_L(Y_{0\bar{M}_1})\} - g\{E_L(Y_{0\bar{M}_0})\}$. Whereas under the second decomposition, the direct and (joint) indirect effects are respectively encoded by: $\gamma_0 = g\{E_L(Y_{1\bar{M}_0})\} - g\{E_L(Y_{1\bar{M}_1})\}$, $\gamma_1 + \gamma_{01} = g\{E_L(Y_{1\bar{M}_1})\} - g\{E_L(Y_{1\bar{M}_0})\}$, so that the direct effect holds the joint distribution of the (counterfactual) mediators fixed under the (hypothetical) control level. Researchers may therefore use such interactions in an effect model to evaluate the practical relevance of different possible decompositions. However, more complex interventional effect models may be required, with additional interaction terms between the covariate(s) and the (mediators') exposure levels when the interventional effects may be modified by baseline covariates. Furthermore, estimation using the inverse weighting procedure may require calculating (non-zero) weights for all individuals that may exacerbate any finite sample biases due to unstable weights. The investigation of weight stabilization strategies such as truncation (Imai and Ratkovic 2015), binning (Naimi, Kaufman, and MacLehose 2014), or optimization-based algorithms (Santacatterina et al. 2019), to improve the inverse weighting estimator should thus be carried out in future work.

Effect heterogeneity in this paper is restricted to pre-specified statistical interactions between baseline covariates and hypothetical exposure levels in the effect model, and thus interaction effects between baseline covariates, and exposure or the mediators, or both, in the fitted mediator and outcome models. Using machine learning-based methods to perform variable selection for covariates that may modify the indirect effects, by extending Imai and Ratkovic (2013) who focus on a single exposure with no mediators, or adapting the debiasing algorithms for heterogeneous treatment effects as in Nie and Wager (2020) or Kennedy (2020), are areas for future work. Nonparametric inference along the lines of the latter two works is additionally promising because it can theoretically eliminate any extrapolation bias arising from an outcome model that is uncongenial with the effect model; such nonparametric inference is achievable, even when parametric effect models are used to provide parsimonious effect summaries (Vansteelandt and Dukes 2020). The proposed methods in this paper are intended to simplify estimation of interventional indirect effects, and is thus relevant to researchers exploring heterogeneous mediated effects via multiple possible mediators, without having to specify a (possibly arbitrary) causal structure among the mediators. When substantive interest is in finding a subset of the (possibly large number of) mediators for further investigation, a threshold that either is pre-determined, such as in Huang and Pan (2016), or controls the family-wise error rate or the false discovery rate under multiple testing scenarios, such as in Dai, Stanford, and LeBlanc (2020) or Derkach et al. (2020), may be employed.

Acknowledgments: The authors would like to thank the Editor, and three reviewers, for their comments on prior versions of this manuscript. Computational resources and services were provided by the VSC (Flemish Supercomputer Center), funded by the FWO and the Flemish Government - department EWI.

Research funding: This research was supported by the Research Foundation - Flanders (FWO) under Grant G019317N.

Author contributions: All authors have accepted responsibility for the entire content of this manuscript and approved its submission. The content is solely the responsibility of the authors and does not represent the official views of the authors' institutions or FWO.

Competing interests: Authors state no conflict of interest.

Informed consent: Not applicable.

References

- Albert, J. M., J. I. Cho, Y. Liu, and S. Nelson. 2019. "Generalized Causal Mediation and Path Analysis: Extensions and Practical Considerations." *Statistical Methods in Medical Research* 28 (6): 1793–807.
- Andrews, R. M., and V. Didelez. 2020. "Insights into the "Cross-World" Independence Assumption of Causal Mediation Analysis." *arXiv Preprint*, arXiv:2003.10341.
- Avin, C., I. Shpitser, and J. Pearl. 2005. "Identifiability of Path-Specific Effects." In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, 357–63. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Dai, J. Y., J. L. Stanford, and M. LeBlanc. 2020. "A Multiple-Testing Procedure for High-Dimensional Mediation Hypotheses." *Journal of the American Statistical Association*: 1–16, <https://doi.org/10.1080/01621459.2020.1765785>.
- Daniel, R. M., B. L. De Stavola, S. N. Cousens, and S. Vansteelandt. 2015. "Causal Mediation Analysis with Multiple Mediators." *Biometrics* 71 (1): 1–14.
- Derkach, A., S. C. Moore, S. M. Boca, and J. N. Sampson. 2020. "Group Testing in Mediation Analysis." *Statistics in Medicine* 39 (18): 2423–36.
- Didelez, V., A. P. Dawid, and S. Geneletti. 2006. "Direct and Indirect Effects of Sequential Treatments." In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, 138–46. Arlington, VA, USA: AUAI Press.
- Efron, B., and R. J. Tibshirani. 1994. *An Introduction to the Bootstrap*. New York, NY: Chapman and Hall/CRC.
- Geneletti, S. 2007. "Identifying Direct and Indirect Effects in a Non-counterfactual Framework." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69: 199–215, <https://doi.org/10.1111/j.1467-9868.2007.00584.x>.
- Greenland, S., J. M. Robins, and J. Pearl. 1999. "Confounding and Collapsibility in Causal Inference." *Statistical Science* 14 (1): 29–46.
- Hong, G. 2010. "Ratio of Mediator Probability Weighting for Estimating Natural Direct and Indirect Effects." In *Proceedings of the American Statistical Association, Biometrics Section*, edited by American Statistical Association, 2401–15. Alexandria, VA, USA.
- Hou, T., W. Dong, R. Zhang, X. Song, F. Zhang, W. Cai, Y. Liu, and G. Deng. 2020. "Self-Efficacy and Fatigue Among Health Care Workers During COVID-19 Outbreak: A Moderated Mediation Model of Posttraumatic Stress Disorder Symptoms and Negative Coping." Preprint (Version 1). Available at Research Square, <https://doi.org/10.21203/rs.3.rs-23066/v1>.
- Huang, Y.-T., and W.-C. Pan. 2016. "Hypothesis Test of Mediation Effect in Causal Mediation Model with High-Dimensional Continuous Mediators." *Biometrics* 72 (2): 402–13.
- Imai, K., and M. Ratkovic. 2013. "Estimating Treatment Effect Heterogeneity in Randomized Program Evaluation." *Annals of Applied Statistics* 7 (1): 443–70.
- Imai, K., and M. Ratkovic. 2015. "Robust Estimation of Inverse Probability Weights for Marginal Structural Models." *Journal of the American Statistical Association* 110 (511): 1013–23.
- Jackson, J. W., and T. J. VanderWeele. 2018. "Decomposition Analysis to Identify Intervention Targets for Reducing Disparities." *Epidemiology* 29 (6): 825–35.
- Kennedy, E. H. 2020. "Optimal Doubly Robust Estimation of Heterogeneous Causal Effects." *arXiv preprint* arXiv:2004.14497.
- Lange, T., M. Rasmussen, and L. C. Thygesen. 2013. "Assessing Natural Direct and Indirect Effects Through Multiple Pathways." *American Journal of Epidemiology* 179 (4): 513–8.
- Lange, T., S. Vansteelandt, and M. Bekaert. 2012. "A Simple Unified Approach for Estimating Natural Direct and Indirect Effects." *American Journal of Epidemiology* 176 (3): 190–5.
- Lin, S.-H., and T. VanderWeele. 2017. "Interventional Approach for Path-Specific Effects." *Journal of Causal Inference* 5 (1), <https://doi.org/10.1515/jci-2015-0027>.
- Lok, J. J. Mar 2019. "Causal Organic Direct and Indirect Effects: Closer to Baron and Kenny." *arXiv Preprint*, art. arXiv:1903.04697.
- Meng, X.-L. 1994. "Multiple-Imputation Inferences with Uncongenial Sources of Input." *Statistical Science* 9 (4): 538–58.
- Micali, N., R. M. Daniel, G. B. Ploubidis, and B. L. De Stavola. 2018. "Maternal Prepregnancy Weight Status and Adolescent Eating Disorder Behaviors: A Longitudinal Study of Risk Pathways." *Epidemiology* 29 (4): 579–89.
- Moreno-Betancur, M., and J. B. Carlin. 2018. "Understanding Interventional Effects: A More Natural Approach to Mediation Analysis?" *Epidemiology* 29 (5): 614–7.
- Moreno-Betancur, M., P. Moran, D. Becker, G. Patton, and J. B. Carlin. July 2020. "Mediation Effects that Emulate a Target Randomised Trial: Simulation-Based Evaluation of Ill-Defined Interventions on Multiple Mediators." *arXiv e-prints*, art. arXiv:1907.06734.

- Naimi, A. I., J. S. Kaufman, and R. F. MacLehose. 2014. "Mediation Misgivings: Ambiguous Clinical and Public Health Interpretations of Natural Direct and Indirect Effects." *International Journal of Epidemiology* 43 (5): 1656–61.
- Nguyen, T. Q., I. Schmid, and E. A. Stuart. Apr 2019. "Clarifying Causal Mediation Analysis for the Applied Researcher: Defining Effects Based on What We Want to Learn." *arXiv Preprint*, art. arXiv:1904.08515.
- Nie, X., and S. Wager. 2020. "Quasi-Oracle Estimation of Heterogeneous Treatment Effects." *arXiv preprint* arXiv:1712.04912.
- Noguchi, Y. 2020. *Pandemic Affects Mental Health of Frontline Health Workers*. Also available at <https://www.npr.org/2020/04/22/841925658/pandemic-affects-mental-health-of-frontline-health-workers>.
- Pearl, J. 2001. "Direct and Indirect Effects." In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, 411–20. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Pearl, J. 2009. *Causality: Models, Reasoning and Inference*, 2nd ed. New York, NY, USA: Cambridge University Press.
- Petersen, M. L., S. E. Sinisi, and M. J. van der Laan. 2006. "Estimation of Direct Causal Effects." *Epidemiology* 17 (3): 276–84.
- Robins, J. M., and S. Greenland. 1992. "Identifiability and Exchangeability for Direct and Indirect Effects." *Epidemiology* 3 (2): 143–55.
- Robins, J. M., and T. S. Richardson. 2010. *Alternative Graphical Causal Models and the Identification of Direct Effects*, 103–58. New York, NY, USA: Oxford University Press. ISBN 9780199754649.
- Robins, J. M. 2000. *Marginal Structural Models versus Structural Nested Models as Tools for Causal Inference*, 95–133. New York, NY, USA: Springer-Verlag.
- Santacatterina, M., C. García-Pareja, R. Bellocco, A. Sönnnerborg, A. M. Ekström, and M. Bottai. 2019. "Optimal Probability Weights for Estimating Causal Effects of Time-Varying Treatments with Marginal Structural Cox Models." *Statistics in Medicine* 38 (10): 1891–902.
- Snowden, J. M., S. Rose, and K. M. Mortimer. 2011. "Implementation of G-Computation on a Simulated Data Set: Demonstration of a Causal Inference Technique." *American Journal of Epidemiology* 173 (7): 731–8.
- Steen, J., T. Loeys, B. Moerkerke, and S. Vansteelandt. 2017. "Flexible Mediation Analysis with Multiple Mediators." *American Journal of Epidemiology* 186 (2): 184–93.
- Taguri, M., J. Featherstone, and J. Cheng. 2018. "Causal Mediation Analysis with Multiple Causally Non-Ordered Mediators." *Statistical Methods in Medical Research* 27 (1): 3–19.
- Tchetgen Tchetgen, E. J. 2014. "A Note on Formulae for Causal Mediation Analysis in an Odds Ratio Context." *Epidemiologic Methods* 2 (1): 21–31.
- VanderWeele, T. J., and E. J. Tchetgen Tchetgen. 2017. "Mediation Analysis with Time Varying Exposures and Mediators." *Journal of the Royal Statistical Society: Series B* 79 (3): 917–38.
- VanderWeele, T. J., and S. Vansteelandt. 2010. "Odds Ratios for Mediation Analysis for a Dichotomous Outcome." *American Journal of Epidemiology* 172 (12): 1339–48.
- VanderWeele, T. J., S. Vansteelandt, and J. M. Robins. 2014. "Effect Decomposition in the Presence of an Exposure-Induced Mediator-Outcome Confounder." *Epidemiology* 25 (2): 300.
- Vansteelandt, S., M. Bekaert, and T. Lange. 2012. "Imputation Strategies for the Estimation of Natural Direct and Indirect Effects." *Epidemiologic Methods* 1 (1): 131–58.
- Vansteelandt, S., and R. M. Daniel. 2017. "Interventional Effects for Mediation Analysis with Multiple Mediators." *Epidemiology* 28 (2): 258–65.
- Vansteelandt, S., and O. Dukes. 2020. "Assumption-Lean Inference for Generalised Linear Model Parameters." *arXiv preprint* arXiv:2006.08402.
- Vansteelandt, S., and N. Keiding. 2011. "Invited Commentary: G-Computation—Lost in Translation?" *American Journal of Epidemiology* 173 (7): 739–42.
- Vansteelandt, S., and T. J. VanderWeele. 2012. "Natural Direct and Indirect Effects on the Exposed: Effect Decomposition under Weaker Assumptions." *Biometrics* 68 (4): 1019–27.

Supplementary Material: The online version of this article offers supplementary material (<https://doi.org/10.1515/em-2020-0023>).